

### In the Claims

1      Claim 1 (currently amended): A method of improving intrusion detection in a computing  
2      network, comprising steps of:

3              defining a plurality of intrusion suspicion levels for use when performing intrusion  
4      detection processing on inbound communications destined for a computing device on the  
5      computing network;

6              for each of a plurality of potential intrusion events, defining a set of at least one condition,  
7      wherein the set describes occurrence of conditions which describe the potential intrusion event;

8              associating one of the defined intrusion suspicion levels with each of the sets, wherein the  
9      associated intrusion suspicion level indicates how suspicious is an inbound communication  
10     matching each condition in the set of conditions;

11              defining a plurality of sensitivity levels for filtering the inbound communications as  
12      potential intrusion events when performing the intrusion detection processing, each of the defined  
13     sensitivity levels usable for a different level of filtering of the inbound communications; and

14              performing the intrusion detection processing for a particular inbound communication  
15      received for the computing device, further comprising steps of:

16              determining whether each condition in any of the at least one sets of conditions are  
17      is matched for the particular inbound communication; and

18              if so, filtering the particular inbound communication by using a currently-applicable  
19      one of the defined sensitivity levels, in concert with the intrusion suspicion level  
20      associated with the set for which each condition is [[the]] matched conditions, to determine if the  
21      particular inbound communication should be treated as an intrusion event.

Claim 2 (canceled)

1       Claim 3 (currently amended): The method according to Claim 1, wherein the determining step  
2       further comprises comparing current conditions in the computing device to [[the]] corresponding  
3       conditions defined in at least one of the sets.

1       Claim 4 (previously presented): The method according to Claim 3, wherein the current conditions  
2       in the computing device comprise contents of the particular inbound communication.

1       Claim 5 (previously presented): The method according to Claim 4, wherein the current conditions  
2       in the computing device further comprise a protocol state of a protocol stack which processes the  
3       particular inbound communication.

1       Claim 6 (previously presented): The method according to Claim 1, further comprising the step of  
2       taking one or more defensive actions upon determining that the particular inbound communication  
3       should be treated as an intrusion event.

1       Claim 7 (original): The method according to Claim 6, wherein the defensive actions are  
2       determined by consulting intrusion detection policy information.

1       Claim 8 (previously presented): The method according to Claim 7, wherein the intrusion

2 detection policy information is stored in a network-accessible repository.

1 Claim 9 (currently amended): The method according to Claim 1, wherein the defined at least one  
2 of the defined sets set of conditions represents one or more an attack signatures signature.

1 Claim 10 (currently amended): The method according to Claim [[9]] 1, wherein at least one of  
2 the defined sets represents attack signatures is a class signature representing for a class of attacks.

1 Claim 11 (currently amended): The method according to Claim 1, wherein each of the at least  
2 one set of conditions is specified as as defined sets comprises a condition part in an intrusion  
3 detection rule, and wherein each of the intrusion detection rules further specifies at least one  
4 action to be taken upon determining that the particular inbound communication should be treated  
5 as an intrusion event.

1 Claim 12 (previously presented): The method according to Claim 1, wherein the performing step  
2 operates in the computing device for which the particular inbound communication is destined.

1 Claim 13 (previously presented): The method according to Claim 12, wherein the performing  
2 step operates within layer-specific intrusion detection logic executing in a protocol stack running  
3 on the computing device.

1 Claim 14 (previously presented): The method according to Claim 1, wherein the performing step

2 operates in a network device which analyzes communications directed to the computing device  
3 for which the particular inbound communication is destined.

1 Claim 15 (currently amended): The method according to Claim 1, wherein the [[using]] filtering  
2 step further comprises consulting using the currently-applicable one of the defined sensitivity  
3 levels and the intrusion suspicion level associated with the matched set to consult a stored  
4 mapping that indicates, for between-each combination of one of the defined sensitivity levels and  
5 [[each]] one of the defined intrusion suspicion levels, whether an inbound communication  
6 corresponding to that combination using the currently-applicable one of the defined sensitivity  
7 levels and the intrusion suspicion level associated with the matched conditions, to determine if the  
8 particular inbound communication should be treated as an intrusion event.

Claims 16 - 21 (canceled)

1 Claim 22 (currently amended): A system for improving intrusion detection in a computing  
2 network, comprising:  
3 a definition of a plurality of intrusion suspicion levels defined for use when performing  
4 intrusion detection processing on inbound communications destined for a computing device on the  
5 computing network;  
6 for each of a plurality of potential intrusion events, a defined definition of a set of at least  
7 one condition, wherein the set describes occurrence of conditions which describe the potential  
8 intrusion event;

9 means for associating one of the defined intrusion suspicion levels with each of the defined  
10 sets, wherein the associated intrusion suspicion level indicates how suspicious is an inbound  
11 communication matching each condition in the set of conditions;

12 a definition of a plurality of sensitivity levels defined for filtering the inbound  
13 communications as potential intrusion events when performing the intrusion detection processing,  
14 each of the defined sensitivity levels usable for a different level of filtering of the inbound  
15 communications; and

16 means for performing the intrusion detection processing for a particular inbound  
17 communication received for the computing device, further comprising:

18 means for determining whether each condition in any of the at least one defined  
19 sets of conditions are is matched for the particular inbound communication; and

20 if so, means for filtering the particular inbound communication by using a  
21 currently-applicable one of the defined sensitivity levels, in concert with the intrusion suspicion  
22 level associated with the set for which each condition is matched conditions, to determine if the  
23 particular inbound communication destined for the computing device should be treated as an  
24 intrusion event.

Claim 23 (canceled)

1 Claim 24 (currently amended): The system according to Claim 22, wherein the means for  
2 determining further comprises means for comparing current conditions in the computing device to  
3 [[the]] corresponding conditions defined in at least one of the sets.

1       Claim 25 (previously presented): The system according to Claim 22, further comprising means  
2       for taking one or more defensive actions upon determining that the particular inbound  
3       communication should be treated as an intrusion event, wherein the defensive actions are  
4       determined by consulting intrusion detection policy information.

1       Claim 26 (currently amended): The system according to Claim 22, wherein each of the ~~at least~~  
2       ~~one set of conditions is specified as defined sets comprises~~ a condition part in an intrusion  
3       detection rule, and wherein each of the intrusion detection rules further specifies at least one  
4       action to be taken upon determining that the particular inbound communication should be treated  
5       as an intrusion event.

1       Claim 27 (currently amended): The system according to Claim 22, wherein the means for  
2       [[using]] filtering further comprises means for using the currently-applicable one of the defined  
3       sensitivity levels and the intrusion suspicion level associated with the matched set to consult  
4       consulting a stored mapping that indicates, for between each combination of one of the defined  
5       sensitivity levels and [[each]] one of the defined intrusion suspicion levels, whether an inbound  
6       communication corresponding to that combination using the currently-applicable one of the  
7       defined sensitivity levels and the intrusion suspicion level associated with the matched conditions;  
8       to determine if the particular inbound communication should be treated as an intrusion event.

Claims 28 - 31 (canceled)

1       Claim 32 (currently amended): A computer program product for improving intrusion detection in  
2       a computing network, the computer program product embodied on one or more computer-  
3       readable media and comprising:

4            computer-readable program code defining a plurality of intrusion suspicion levels for use  
5       when performing intrusion detection processing on inbound communications destined for a  
6       computing device on the computing network;

7            for each of a plurality of potential intrusion events, computer-readable program code  
8       defining a set of at least one condition, wherein the set describes occurrence of conditions which  
9       describe the potential intrusion event;

10          computer-readable program code associating one of the defined intrusion suspicion levels  
11       with each of the sets, wherein the associated intrusion suspicion level indicates how suspicious is  
12       an inbound communication matching each condition in the set of conditions;

13          computer-readable program code defining a plurality of sensitivity levels for filtering the  
14       inbound communications as potential intrusion events when performing the intrusion detection  
15       processing, each of the defined sensitivity levels usable for a different level of filtering of the  
16       inbound communications; and

17          computer-readable program code for performing the intrusion detection processing for a  
18       particular inbound communication received for the computing device, further comprising:

19            computer-readable program code for determining whether each condition in any of  
20       the at least one sets of conditions are is matched for the particular inbound communication; and

21            if so, computer-readable program code for filtering the particular inbound

22 communication by using a currently-applicable one of the defined sensitivity levels, in concert with  
23 the intrusion suspicion level associated with the matched-conditions set, to determine if the  
24 particular inbound communication should be treated as an intrusion event.

Claim 33 (canceled)

1 Claim 34 (currently amended): The computer program product according to Claim 32, wherein  
2 the computer-readable program code for determining further comprises computer-readable  
3 program code for comparing current conditions in the computing device to [[the]] corresponding  
4 conditions defined in at least one of the sets, the current conditions in the computing device  
5 comprising contents of the particular inbound communication.

1 Claim 35 (currently amended): The computer program product according to Claim 32, wherein  
2 the computer-readable program code for determining further comprises computer-readable  
3 program code for comparing current conditions in the computing device to [[the]] corresponding  
4 conditions defined in at least one of the sets, the current conditions in the computing device  
5 comprising contents of the particular inbound communication and a protocol state of a protocol  
6 stack which processes the particular inbound communication.

1 Claim 36 (previously presented): The computer program product according to Claim 32, further  
2 comprising computer-readable program code for taking one or more defensive actions upon  
3 determining that the particular inbound communication should be treated as an intrusion event,

4 wherein the defensive actions are determined by consulting intrusion detection policy information  
5 stored in a policy repository.

1 Claim 37 (currently amended): The computer program product according to Claim 32, wherein  
2 the defined at least one of the defined sets set of conditions represents one or more an attack  
3 signatures signature, and wherein at least one of the attack signatures is a class signature  
4 representing a class of attacks.

1 Claim 38 (previously presented): The computer program product according to Claim 32, wherein  
2 the computer-readable program code for performing operates in the computing device for which  
3 the particular inbound communication is destined.

1 Claim 39 (previously presented): The computer program product according to Claim 32, wherein  
2 the computer-readable program code for performing operates in a network device which analyzes  
3 communications directed to the computing device for which the particular inbound  
4 communication is destined.

1 Claim 40 (currently amended): The computer program product according to Claim 32, wherein  
2 the computer-readable program code for [[using]] filtering further comprises computer-readable  
3 code for consulting using the currently-applicable one of the defined sensitivity levels and the  
4 intrusion suspicion level associated with the matched set to consult a stored mapping that  
5 indicates, for between each combination of one of the defined sensitivity levels and [[each]] one of

6 the defined intrusion suspicion levels, whether an inbound communication corresponding to that  
7 combination using the currently-applicable one of the defined sensitivity levels and the intrusion  
8 suspicion level associated with the matched conditions set, to determine if the particular inbound  
9 communication should be treated as an intrusion event.

Claims 41 - 44 (canceled)

1 Claim 45 (currently amended): The method according to Claim 6, wherein the defensive actions  
2 are specified as actions in a rule in which the matched-conditions are set is specified.

1 Claim 46 (previously presented): The method according to Claim 6, wherein at least one of the  
2 defensive actions comprises discarding the particular inbound communication.

1 Claim 47 (previously presented): The method according to Claim 6, wherein at least one of the  
2 defensive actions comprises limiting at least one of resources or traffic associated with a  
3 connection on which the particular inbound communication is received.

1 Claim 48 (previously presented): The method according to Claim 6, wherein at least one of the  
2 defensive actions comprises dynamically dropping a deny filter into the computing network to  
3 shun subsequent traffic.

1 Claim 49 (previously presented): The method according to Claim 6, wherein at least one of the

2 defensive actions comprises reporting the intrusion event to one or more entities.

1 Claim 50 (previously presented): The method according to Claim 49, wherein reporting the  
2 intrusion event to one or more entities further comprises sending an alert to a management  
3 component external from the computing device for which the particular inbound communication is  
4 destined.

1 Claim 51 (previously presented): The method according to Claim 49, wherein reporting the  
2 intrusion event to one or more entities further comprises writing at least one event record to at  
3 least one of a system log and a console.

1 Claim 52 (previously presented): The method according to Claim 49, wherein reporting the  
2 intrusion event to one or more entities further comprises recording inbound communications  
3 associated with the intrusion event in at least one of a trace or other repository.

1 Claim 53 (previously presented): The method according to Claim 49, wherein reporting the  
2 intrusion event to one or more entities further comprises writing statistics records on normal  
3 behavior to establish baselines as to what constitutes abnormal behavior for the inbound  
4 communications.

1 Claim 54 (currently amended): The method according to Claim 1, wherein at least one of the  
2 defined sets of conditions specifies a current system state of the computing device.

1       Claim 55 (currently amended): The method according to Claim 1, wherein at least one of the  
2       defined sets-of-conditions specifies at least one threshold reached at the computing device.

1       Claim 56 (currently amended): The method according to Claim 1, wherein at least one of the  
2       defined sets-of-conditions specifies at least one state transition to be caused, at the computing  
3       device, upon receiving the particular inbound communication.

1       Claim 57 (previously presented): The method according to Claim 1, wherein the currently-  
2       applicable sensitivity level is specified, for the computing device, by a systems administrator.

1       Claim 58 (previously presented): The method according to Claim 1, wherein the currently-  
2       applicable sensitivity level is specified, for the computing device, by configuration data in a stored  
3       repository.